

ATTY. DOCKET NO. SAMS01-00269

PATENT

APPARATUS AND METHOD FOR PERFORMING HIGH-SPEED
LOOKUPS IN A ROUTING TABLE

Inventor(s) :

Jack C. Wybenga
2129 Stone Creek
Plano
Collin County
Texas 75075
United States citizen

Patricia K. Sturm
4424 Rushing Road
Dallas
Collin County
Texas 75287
United States citizen

Assignee:

SAMSUNG ELECTRONICS Co., LTD.
416, Maetan-dong, Paldal-gu
Suwon-city, Kyungki-do
Republic of Korea

William A. Munck
John T. Mockler
Davis Munck, P.C.
Three Galleria Tower
13155 Noel Road, Suite 900
Dallas, Texas 75240
(972) 628-3600

APPARATUS AND METHOD FOR PERFORMING HIGH-SPEED
LOOKUPS IN A ROUTING TABLE

5

TECHNICAL FIELD OF THE INVENTION

[001] The present invention relates to massively parallel routers and, more specifically, to a massively parallel, distributed architecture router that contains a routing (or forwarding) lookup mechanism capable of performing high-speed lookups.

BACKGROUND OF THE INVENTION

[002] There has been explosive growth in Internet traffic due to the increased number of Internet users, various service demands from those users, the implementation of new services, such as voice-over-IP (VoIP) or streaming applications, and the development of mobile Internet. Conventional routers, which act as relaying nodes connected to sub-networks or other routers, have accomplished their roles well, in situations in which the time required to process packets, determine their destinations, and forward the packets to the destinations is usually smaller than the transmission time on network paths. More recently, however, the packet transmission capabilities of high-bandwidth network paths

and the increases in Internet traffic have combined to outpace the processing capacities of conventional routers. Increasingly, routers are the cause of major bottlenecks in the Internet.

[003] The limitations of conventional routers have been at least partially overcome by the introduction of massively parallel, distributed architecture routers. The use of optical connections has also greatly increased throughput. However, even massively parallel, distributed architecture routers have problems caused in part by the use of routing tables (or forwarding tables) that perform address translation lookups, among other things. Line speeds are increasing faster than processing speeds that perform routing table lookups. Since route lookups require the longest prefix match, this is a non-trivial problem. Internet Protocol Version 6 (IPv6) has aggravated this problem, because IPv6 uses 128-bit addresses, compared to the 32-bit addresses used in IPv4. Adding Type of Service (TOS) and Layer 4 addressing fields into the lookup value makes the problem still worse.

[004] Routing tables with a million entries are not uncommon. Some lookup schemes (e.g., hashing, digital trees) are able to reduce the search time, but the large number of routing table entries leads to memory problems. It is prohibitively expensive and technically difficult to incorporate enough high-speed memory

to support fairly flat search tables. As the number of memory chips increases, maintaining high performance becomes difficult due to layout considerations. As a result, memory access times are too slow to permit very deep search tables.

5 [005] Some proposals use ternary content addressable memory (TCAM) devices to increase lookup speeds, but these devices are impractical due to expense and power consumption. Placing enough TCAMs on a circuit card to handle forwarding table lookups up to 144 bits wide, with up to a million entries, is prohibitive in both
10 respects. Hashing and state-of-the-art search techniques are not adequate to keep routing tables within a reasonable size for cost and performance considerations and to keep the number of lookup stages low enough that memory access times enable lookups to keep up with line speeds. Thus, there is no practical method for doing
15 IPv6 lookups at line speed for high-speed interfaces.

[006] Therefore, there is a need in the art for an improved high-speed router. In particular, there is a need for an improved routing (forwarding) lookup mechanism that can perform lookups at line speed for a high-speed interface.

SUMMARY OF THE INVENTION

[007] The present invention provides an apparatus for performing IPv4 and IPv6 routing table lookups at line speeds of 10 gigabits per second (Gbps) and higher. IPv4 lookups of 50 bits enable forwarding based on 32 bits of Layer 3 IP address, 12 bits of Layer 4 address (sockets), and 6 bits of TOS. IPv6 lookups of 144 bits enable forwarding based on 128 bits of Layer 3 address and up to 16 bits of Layer 4 addressing and TOS.

[008] The present invention uses a combination of hashing, digital search trees, and pipelining to achieve the goals of line speed routing table lookup operations in lookup tables containing up to one million entries. Advantageously, the present invention may be implemented with relatively low cost parts. A key aspect of the present invention is the use of a trie-based scheme to keep the size of the lookup structures within the practical limits of high-speed SRAM. Only the final stage of the route lookup must reside in low-speed DRAM. A pipelined hardware lookup scheme achieves the high throughput.

[009] To address the above-discussed deficiencies of the prior art, it is a primary object of the present invention to provide, for use in a router, a lookup circuit for translating received addresses into destination addresses. According to an advantageous

embodiment, the lookup circuit comprises M pipelined memory circuits for storing a trie table capable of translating a first received address into a first destination address. The M memory circuits are pipelined such that a first portion of the first
5 received address accesses an address table in a first memory circuit and an output of the first memory circuit accesses an address table in a second memory circuit.

[010] According to one embodiment of the present invention, the output of the first memory circuit comprises a first address
10 pointer that indexes a start of the address table in the second memory circuit.

[011] According to another embodiment of the present invention, the first address pointer and a second portion of the first received address access the address table in the second memory
15 circuit.

[012] According to still another embodiment of the present invention, an output of the second memory circuit accesses an address table in a third memory circuit.

[013] According to yet another embodiment of the present
20 invention, the output of the second memory circuit comprises a second address pointer that indexes a start of the address table in the third memory circuit.

[014] According to a further embodiment of the present invention, address pointers output from the M pipelined memory circuits are selectively applied to a final memory circuit storing a routing table, the routing table comprising a plurality of
5 destination addresses associated with the received addresses.

[015] According to a still further embodiment of the present invention, the lookup circuit further comprises a memory interface capable of selectively applying to the final memory circuit an address pointer associated with the first received address and an
10 address pointer associated with a subsequently received address, such that the address pointer associated with the first received address is applied to the final memory circuit prior to the address pointer associated with the subsequently received address.

[016] According to a yet further embodiment of the present
15 invention, the M pipelined memory circuits comprise static random access memory (SRAM) circuits and the final memory circuit comprises a dynamic random access memory (DRAM) circuit.

[017] Before undertaking the DETAILED DESCRIPTION OF THE INVENTION below, it may be advantageous to set forth definitions of
20 certain words and phrases used throughout this patent document: the terms "include" and "comprise," as well as derivatives thereof, mean inclusion without limitation; the term "or," is inclusive,

meaning and/or; the phrases "associated with" and "associated therewith," as well as derivatives thereof, may mean to include, be included within, interconnect with, contain, be contained within, connect to or with, couple to or with, be communicable with, cooperate with, interleave, juxtapose, be proximate to, be bound to or with, have, have a property of, or the like; and the term "controller" means any device, system or part thereof that controls at least one operation, such a device may be implemented in hardware, firmware or software, or some combination of at least two of the same. It should be noted that the functionality associated with any particular controller may be centralized or distributed, whether locally or remotely. Definitions for certain words and phrases are provided throughout this patent document, those of ordinary skill in the art should understand that in many, if not most instances, such definitions apply to prior, as well as future uses of such defined words and phrases.

BRIEF DESCRIPTION OF THE DRAWINGS

[018] For a more complete understanding of the present invention and its advantages, reference is now made to the following description taken in conjunction with the accompanying
5 drawings, in which like reference numerals represent like parts:

[019] FIGURE 1 illustrates a distributed architecture router that implements a fast lookup forwarding table according to the principles of the present invention;

[020] FIGURE 2 illustrates selected portions of an exemplary
10 routing node in the distributed architecture router in FIGURE 1 according to one embodiment of the present invention; and

[021] FIGURE 3 illustrates a trie-based, pipelined routing table according to the principles of the present invention; and

[022] FIGURE 4 is a timing diagram illustrating the operation
15 of the trie-based pipelined routing table in FIGURE 3.

DETAILED DESCRIPTION OF THE INVENTION

[023] FIGURES 1 through 4, discussed below, and the various embodiments used to describe the principles of the present invention in this patent document are by way of illustration only and should not be construed in any way to limit the scope of the invention. Those skilled in the art will understand that the principles of the present invention may be implemented in any suitably arranged distributed router.

[024] FIGURE 1 illustrates exemplary distributed architecture router 100, which implements a fast lookup forwarding table according to the principles of the present invention. Distributed architecture router 100 provides scalability and high-performance using up to N independent routing nodes (RN), including exemplary routing nodes 110, 120, 130 and 140, connected by switch 150, which comprises a pair of high-speed switch fabrics 155a and 155b. Each routing node comprises an input-output processor (IOP) module, and one or more physical medium device (PMD) module. Exemplary RN 110 comprises PMD module 112 (labeled PMD-a), PMD module 114 (labeled PMD-b), and IOP module 116. RN 120 comprises PMD module 122 (labeled PMD-a), PMD module 124 (labeled PMD-b), and IOP module 126. RN 130 comprises PMD module 132 (labeled PMD-a), PMD module 134 (labeled PMD-b), and IOP module 136. Finally, exemplary RN 140

comprises PMD module 142 (labeled PMD-a), PMD module 144 (labeled PMD-b), and IOP module 146.

[025] Each one of IOP modules 116, 126, 136 and 146 buffers incoming Internet protocol (IP) frames and MPLS frames from subnets or adjacent routers, such as router 190 and network 195. Additionally, each of IOP modules 116, 126, 136 and 146 classifies requested services, looks up destination addresses from frame headers or data fields, and forwards frames to the outbound IOP module. Moreover, each IOP module also maintains an internal routing table determined from routing protocol messages and provisioned static routes and computes the optimal data paths from the routing table. Each IOP module processes an incoming frame from one of its PMD modules. According to one embodiment of the present invention, each PMD module encapsulates an incoming frame (or cell) from an IP network (or ATM switch) for processing in an IOP module and performs bus conversion functions.

[026] Each one of routing nodes 110, 120, 130, and 140, configured with an IOP module and PMD module(s) and linked by switch fabrics 155a and 155b, is essentially equivalent to a router by itself. Thus, distributed architecture router 100 can be considered a set of RN building blocks with high-speed links (i.e., switch fabrics 155a and 155b) connected to each block. Switch

fabrics 155a and 155b support frame switching between IOP modules.

Switch processor (SWP) 160a and switch processor (SWP) 160b, located in switch fabrics 155a and 155b, respectively, support system management.

5 [027] Unlike a traditional router, distributed architecture router 100 requires an efficient mechanism of monitoring the activity (or "aliveness") of each routing node 110, 120, 130, and 140. Distributed architecture router 100 implements a routing coordination protocol (called "loosely-coupled unified environment
10 (LUE) protocol") that enables all of the independent routing nodes to act as a single router by maintaining a consistent link-state database for each routing node. The loosely-unified environment (LUE) protocol is based on the design concept of OSPF (Open Shortest Path First) routing protocol and is executed in parallel
15 by daemons in each one of RN 110, 120, 130, and 140 and in SWP 160a and SWP 160b to distribute and synchronize routing tables. As is well known, a daemon is an agent program that continuously operates on a processing node and provides resources to client systems. Daemons are background processes used as utility functions.

20 [028] FIGURE 2 illustrates selected portions of exemplary routing node 120 in distributed architecture router 100 according to one embodiment of the present invention. Router 100 shares

routing information in the form of aggregated routes among the routing engines. The routing engines are interconnected through Gigabit optical links to the switch modules (SWMs). Multiple SWMs can be interconnected through 10 Gbps links. Classification module 230 is an optional daughter card that may be inserted on any or all IOP modules. Ingress data can be sent to classification modules 230 to enable, for example, IPv6 tunneling through router 100, streams-based billing, subnet independent NAT, Layers 4-7 and QoS-based forwarding, data filtering and blocking for firewall functionality, and data surveillance, among other functions.

[029] Routing node 120 comprises physical medium device (PMD) module 122, physical medium device (PMD) module 124 and input-output processor module 126. PMD module 122 (labeled PMD-a) comprises physical layer circuitry 211, physical medium device (PMD) processor 213 (e.g., IXP 1240 processor), and peripheral component interconnect (PCI) bridge 212. PMD module 124 (labeled PMD-b) comprises physical layer circuitry 221, physical medium device (PMD) processor 223 (e.g., IXP 1240 processor), and peripheral component interconnect (PCI) bridge 222.

[030] IOP module 126 comprises classification module 230, system processor 240 (e.g., MPC 8245 processor), network processor 260 (e.g., IXP 1200 or IXP 1240 processor), peripheral component

interconnect (PCI) bridge 270, and Gigabit Ethernet connector 280.

Classification module 230 comprises content addressable memory (CAM) 231, classification processor 232 (e.g., MPC 8245 processor), classification engine 233 and custom logic array (CLA) 234 (e.g.,
5 FPGA). Classification engine 233 is a state graph processor. Custom logic array 234 controls the flow of the packet within classification module 230 and between classification module 230 and network processor 260. PCI bus 290 connects PCI bridges 212, 222 and 270, classification processor 232, and system processor 240 for
10 control plane data exchange such as route distribution. IX bus 296 interconnects PMD processor 213, PMD processor 223, and network processor 260 for data plane traffic flow. Local bus 292 interconnects classification module 230 and network processor 260 for data plane traffic flow.

15 **[031]** Network processor 260 comprises microengines that perform frame forwarding and a control plane processor. Network processor 260 uses distributed forwarding table (DFT) 261 to perform forwarding table lookup operations. The network processor (e.g., network processor 260) in each IOP module (e.g., IOP module 126)
20 performs frame forwarding using a distributed forwarding table (e.g., DFT 261).

[032] As the foregoing description illustrates, router 100 contains a number of routing (forwarding) tables that translate IPv4 and IPv6 prefixes into destination addresses. As the line speeds of router 100 increase to the 10 gigabit per second (Gbps) range, such as in an OC-192c optical link, the lookup speeds of the routing tables are required to be very fast. The lookup speed is limited in part by the length of the longest matching prefix of an IPv4 or and IPv6 address.

[033] A number of approaches have been used to search for the longest matching prefixes. Most approaches use one of two methods: 1) a search tree method; or a 2) search trie method. A search tree checks the value of the entry with the median value of each sub-tree. If the value is less than the median value, it is directed to the left half of the sub-tree and if it is larger, it is pointed to the right half.

[034] A search trie uses a "thumb indexing" method, as in a dictionary. Each bit in the address is checked and a Logic 0 points to the left half of the sub-tree and a Logic 1 points to the right half of the subtree. The trie is traversed until a leaf node is reached which determines the longest matching prefix. In the worst case, the number of memory accesses required for these

schemes to determine the longest matching prefix equals the depth,
D given by:

$$D = (\text{Address Bits}) / \log_2(M), \quad [\text{Eqn. 1}]$$

where M is the degree of the trie (i.e., the number of ways to
5 branch at each stage of the lookup) and $\log_2(M)$ is the number of
bits consumed in each stage of the lookup. Most trie-based schemes
attempt to reduce the number of memory accesses by reducing the
trie depth.

[035] Router 100 meets the requirements imposed by high line
10 speeds by implementing a trie-based memory architecture that
includes pipelined memory stages that perform very fast lookups. A
final stage is a dynamic random access memory (DRAM) circuit that
contains the routing table entries. The preceding pipeline stages
are made from very fast static random access memory (SRAM) circuits
15 that contain address pointers that index into subsequent pipeline
stages or into the routing tables entries in the final DRAM stage.

[036] The expected SRAM memory (bits/entry) of a trie for n
random uniformly distributed entries is given by:

$$E(\text{Mem}(\text{Bits/Entry})) = M / \ln(M), \quad [\text{Eqn. 2}]$$

20 where M is the degree of the trie structure.

[037] It is possible to calculate the maximum SRAM requirement
and the expected SRAM requirement for different degrees of the

trie. The maximum SRAM requirement arises from extreme cases that generally are not observed in conventional routing tables. It is further noted that the computed expected SRAM is less than that required for the actual routing tables. Therefore, the expected
5 SRAM required is calculated and a scaling factor is used to allocate sufficient SRAM for the desired routing table size. The depth of the trie is also dependent on the degree of the trie. The depth of the trie determines the number of SRAM accesses.

[038] The expected SRAM for a 64 degree trie is 15.39 bits per
10 entry from Equation 2. This is approximately 16 megabits for a one million entry table. Using a scaling factor of 5 to provide sufficient space for actual IPv6 routing tables gives an SRAM requirement of approximately 80 megabits. Practical memory performance considerations for laying out circuit cards with 20
15 nanosecond memory chips give an expected SRAM limit of about 200 megabits. Two copies of the trie tables are maintained to allow seamless table updates. Thus, it is noted that a trie of degree 64 is near the 200 Mbit SRAM limit. This is a good indication that the degree of the trie structure should be no more than 64.

20 [039] Suppose, as an example, that a 16 degree trie is proposed. The memory requirements to store a one million entry table are:

SRAM = 5 x 6 Mbit = 30 Mbit; and

DRAM = 8 x 30 Mbit = 240 Mbit.

[040] Equation 2 gives 5.77 bits per entry for a 16 degree trie, hence approximately 6 Mbits for a million entry table. A scaling factor of 5 is used to provide sufficient space for actual IPv6 routing tables. For IPv4 routing tables, a scaling factor of 3 could be used. By assuming 8 bits for storing port numbers and noting that each trie entry may be a leaf, the DRAM requirement of 240 megabits is found.

10 [041] The present invention starts by hashing a fixed number of bits to find the starting point for the trie table search for longest prefix match. The packet is classified using header information, such as type of service. The classification information and high order destination address bits are used for this hashing function. When classification bits are used, the length of the search increases and more stages of the lookup mechanism may be required.

[042] According to an exemplary embodiment of the present invention, no classification is done and it is assumed that the IPv4 and IPv6 prefixes seen by the router are never shorter than 16 bits, so the first 16 bits can be hashed to provide the starting point for the trie table search for longest prefix match. Thus,

the trie lookup is done on the remaining 16 bits for a 32-bit IPv4 address and on the remaining 112 bits for a 128-bit IPv6 address. From Equation 1, the depth of this 16 degree trie for IPv4 lookups is 4 and for IPv6 lookups is 28. If classification were done, more
5 bits would be used and the depth of the IPv4 and IPv6 lookups would be greater.

[043] With a minimum data packet size of 64 bytes and a usable throughput of 76% of bus bandwidth, a 1 Gigabit per second (Gbps) Ethernet interface can support approximately 1.5 million data
10 packets per second. A rate of 1.5 million lookups per second corresponds to 666 nanoseconds for each data packet. Therefore, the time available for each SRAM level is equal to:

$$(666 \text{ ns}) / (28 \text{ levels}) = 23.8 \text{ ns/level}$$

for the worst case IPv6 lookups. With an SRAM cycle time of 8
15 nanoseconds, each level can easily be searched in 23.8 ns. Such an implementation can even be done using off the shelf FPGA and SRAM chips. No pipelining is necessary for such an implementation. Large sizes can be implemented easily.

[044] However, for a 10 Gbps Ethernet interface, the number of
20 lookups per second increases to 15 million, leaving only $67/28 = 2.4$ nanoseconds per lookup. This is not achievable with current SRAM circuits. Thus, a new approach is needed. The present

invention provides this new approach by pipelining the lookup stages.

[045] The maximum lookup rate that can be achieved by using a single RLDRAM is 40 million lookups per second or 25 nanoseconds per lookup. To get maximum throughput from such a scheme, one embodiment of the present invention may employ a 16 degree trie that has a depth of 4 levels for IPv4 and memory requirements of:

SRAM = 3 x 6 Mbit = 18 Mbit; and

DRAM = 8 x 18 Mbit = 144 Mbit.

10 [046] A scaling factor of 3 is used for SRAM, which retains a reasonable SRAM size. For IPv6, a scaling factor of 5 is assumed.

The scheme could easily be implemented using a four-stage on-chip pipeline and a final RLDRAM pipeline stage, as shown in FIGURE 3.

[047] FIGURE 3 illustrates a trie-based, pipelined routing table 300 according to the principles of the present invention. In one embodiment, routing table 300 may represent distributed forwarding table 261, for example. Four bits are consumed in each memory access in each pipeline stage. It is noted that more stages of pipelining are needed for IPv6. Routing table 300 comprises address buffer 310, static random access memory (SRAM) circuits 321-325, memory interface 330, memory controller 340, and dynamic random access memory (DRAM) circuit 350. SRAM circuits 321-325

15
20

contain address pointers that index into subsequent ones of the SRAM circuits in the pipeline or into the routing table entries in DRAM circuit 350.

[048] Memory controller 340 controls the storing of addresses
5 into address buffer 310 and controls the selective outputting of portions of each address in address buffer 310 to each one of SRAM circuits 321-325. Memory controller 340 also controls the storing of the outputs of SRAM circuit 321-325 into memory interface 330 and controls the selective outputting of addresses in memory
10 interface 330 to DRAM circuit 350.

[049] Address buffer 310 receives and buffers 32-bit IPv4 addresses. Each 32-bit address, A[31:0] is logically divided into a first 16-bit portion, A[31:16], and four other 4-bit portions, A[15:12], A[11:8], A[7:4], and A[3:0]. The address portions are
15 applied to SRAM circuits 321-325 over five (5) sequential time slots. Address bits A[31:16] are applied to SRAM circuit 321 during a first time slot. The output of SRAM circuit 321 and address bits A[15:12] are applied to SRAM circuit 322 during a second time slot. The output of SRAM circuit 322 and address bits
20 A[11:8] are applied to SRAM circuit 323 during a third time slot. The output of SRAM circuit 323 and address bits A[7:4] are applied to SRAM circuit 324 during a fourth time slot. The output of SRAM

circuit 324 and address bits A[3:0] are applied to SRAM circuit 325 during a fifth time slot.

[050] All subsequent addresses are applied in a similar manner to SRAM circuits 321-325. The subsequent addresses are also
5 applied from address buffer 310 in a pipelined manner. Thus, during the second time slot, when address bits A[15:12] of a first sequential 32-bit address are being applied to SRAM circuit 322, the address bits A[31:16] of a second sequential 32-bit address are being applied to SRAM circuit 321. Then, during the third time
10 slot, address bits A[11:8] of the first sequential 32-bit address are applied to SRAM circuit 323 at the same time that address bits A[15:12] of the second sequential 32-bit address are applied to SRAM circuit 322 and address bits A[31:16] of a third sequential 32-bit address are applied to SRAM circuit 321.

15 [051] As noted above, it is assumed that router 100 does not see IPv4 prefixes shorter than 16 bits. Thus, the first sixteen address bits, A[31:16], are applied together to SRAM circuit 321. SRAM circuit 321 contains a table of address pointers having 2^{16} entries (i.e., 64K entry table). Each table entry contains an
20 address pointer and a flag bit indicating whether or not the address translation is complete. If the IPv4 prefix is only 16 bits long, then the flag bit is set and the address pointer is

latched into memory interface 330 in order to be applied to DRAM 350. If the IPv4 prefix is longer than 16 bits, then the flag is not set and the address pointer from SRAM circuit 321 is applied to SRAM circuit 322.

5 [052] If smaller prefixes were seen (e.g., 8 bit prefixes), then the size of initial SRAM table 321 would decrease to 2^8 entries and the number of stages in the lookup mechanism would increase by two. Use of classification bits could increase the size of SRAM 321 or increase the number of lookup stages.

10 [053] SRAM circuit 322 contains a maximum of N tables, where N is determined by the size of the table in SRAM circuit 321. Each of the N tables in SRAM circuit 322 contains 16 entries. The start of each table is indexed by the address pointer from SRAM circuit 321. Address bits A[15:12] are used to select a particular one of
15 the 16 entries in the table indexed by the address pointer from SRAM circuit 321.

 [054] Each of SRAM circuits 322-325 contains N tables that operate in a similar to the table of address pointers in SRAM circuit 321. For example, in SRAM circuit 322, each table entry
20 contains an address pointer and a flag bit indicating whether or not the address translation is complete. If the IPv4 prefix is 20 bits long, then the flag bit is set and the address pointer from

SRAM circuit 322 is latched into memory interface 330 in order to be applied to DRAM 350. If the IPv4 prefix is longer than 20 bits, then the flag is not set and the address pointer from SRAM circuit 322 is applied to SRAM circuit 323. This process continues through
5 SRAM circuits 323, 324 and 325. Memory controller 340 detects when a flag bit is set after each SRAM stage and controls the latching of the address pointer into memory interface 330.

[055] FIGURE 4 is a timing diagram illustrating the operation of the trie-based pipelined routing table in FIGURE 3. Three
10 addresses, Address 1, Address 2, and Address 3, are propagated through the pipeline stages of routing table 300. In an exemplary embodiment, it is assumed that SRAM circuits 321-325 and DRAM circuit 350 each have a 25 nanosecond propagation time. Thus, times T0-T8 are each spaced 25 nanoseconds apart. In FIGURE 4,
15 each black square indicates that a prefix match has occurred and that the output from the SRAM circuit is the final address in the routing tables in DRAM circuit 350.

[056] At time T0=0, Address A1 is applied SRAM circuit 321 and at time T1 = 25 nanoseconds, an address pointer is output by SRAM
20 circuit 321. The empty square indicates that a prefix match has not occurred (flag not set) for Address 1 and the address pointer from SRAM circuit 321 is used as an index into SRAM circuit 322.

At time T2 = 50 nanoseconds, an address pointer is output by SRAM circuit 322. The empty square indicates that a prefix match has not occurred for Address 1 and the address pointer from SRAM circuit 322 is used as an index into SRAM circuit 323. At time T3 = 75 nanoseconds, an address pointer is output by SRAM circuit 323.

The empty square indicates that a prefix match has not occurred for Address 1 and the address pointer from SRAM circuit 323 is used as an index into SRAM circuit 324. At time T4 = 100 nanoseconds, an address pointer is output by SRAM circuit 324. The empty square indicates that a prefix match has not occurred for Address 1 and the address pointer from SRAM circuit 324 is used as an index into SRAM circuit 325.

[057] Finally, at time T5 = 125 nanoseconds, an address pointer is output by SRAM circuit 325. The black square indicates that a prefix match has occurred for Address 1 and the address pointer from SRAM circuit 325 is used as an index into DRAM circuit 350. Memory controller 340 detects that the flag from SRAM circuit 325 is set and causes memory interface to transfer the address pointer to DRAM circuit 350. At time T6 = 150 nanoseconds, DRAM circuit 350 outputs a destination address, indicated by a square containing an "X". It is assumed in the example that the delay time of memory interface 330 is negligibly small so that the delays of memory

interface 330 and DRAM circuit 350 are collectively shown as 25 nanoseconds.

[058] A similar process occurs for Address 2, except that Address 2 trails Address 1 by one SRAM stage (i.e., 25 nanoseconds) and a prefix match occurs at time T4, when SRAM circuit 323 outputs an address pointer with the flag set. Memory controller 340 detects that the flag from SRAM circuit 323 is set and causes memory interface 330 to retain the address pointer from SRAM circuit 323. It is noted that the Address 2 match occurs before the Address 1 match. However, memory controller 340 and memory interface 330 maintain the order of the address pointers for Address 1 and Address 2. Thus, memory interface 330 applies the address pointer from SRAM circuit 325 to DRAM circuit 350 at time T5 and applies the address pointer from SRAM circuit 323 to DRAM circuit 350 at time T6, one time slot (i.e., 25 nanoseconds) after time T5. DRAM circuit 350 outputs the destination address for Address 2 at time T7, indicated by a box containing an "X".

[059] A similar process occurs for Address 3, except that Address 3 trails Address 2 by one SRAM stage (i.e., 25 nanoseconds) and a prefix match occurs at time T6, when SRAM circuit 324 outputs an address pointer with the flag set. Again, however memory controller 340 and memory interface 330 maintain the order to the

address pointers for Address 3 and Address 2. Thus, memory interface 330 applies the address pointer from SRAM circuit 323 to DRAM circuit 350 at time T6 and applies the address pointer from SRAM circuit 324 to DRAM circuit 350 at time T7, one time slot
5 (i.e., 25 nanoseconds) after time T6. DRAM circuit 350 outputs the destination address for Address 3 at time T8, indicated by a box containing an "X".

[060] Since destination addresses emerge from DRAM circuit 350 every 25 nanoseconds, routing table 300 is capable of 40 million
10 lookups per second. An OC-192c optical link requires 24 million lookups per second for IPv4, assuming 40 byte packets and subtracting the OC-192c and packet framing overhead. This will reduce for IPv6 due to larger minimum size packets. The extra time for IPv6 may be utilized for giving more time at each SRAM level.
15 Thus, the present invention is very scalable.

[061] When routing table 300 is updated, the following actions are necessary:

i) Port Reassignment - If an existing prefix is simply reassigned to a different output port, then a single DRAM write is
20 required for each change. No changes are needed for the SRAM tables;

ii) New prefix insertion - Whenever a new prefix is inserted into the lookup table, the tables are rebuilt from scratch. The time required to do this is limited primarily by the time to write the table to DRAM circuit 350. For a one million entry table, the
5 time required is about four (4) milliseconds. If this dead time is unacceptable, then two copies of DRAM circuit 350 and possibly SRAM circuit 321-325 may be needed; and

iii) Table Calculation - It is estimated to take about 100 milliseconds to calculate the routing (forwarding) table using a
10 250 MIPS processor.

[062] The present invention is capable of providing line speed lookups for 1 Gbps and 10 Gbps interfaces. There is a trade-off between memory size (especially SRAM trie table storage) and the number of lookups that must be done for each data packet (i.e.
15 trie depth). As noted above, the performance of SRAM circuits limits the amount of SRAM to about 200 Mbits. Due to relatively long trie table update times, two copies of the trie tables may be required - one to perform searches (lookups) while the other is updated. This limits the amount of SRAM available for each trie
20 table to about 100 Mbits.

[063] Additionally, SRAM and DRAM lookup rates limit the trie depth to about 32 stages. Thus, a degree 16 trie table is

advantageous. With a requirement for one million forwarding table entries, a 1 Gbps Ethernet interface may be serviced by a degree 16 trie table, 30 Mbits of SRAM, and 240 Mbits of DRAM for IPv6. Achieving line rate lookups for 10 Gbps Ethernet or OC-192c
5 interfaces requires hardware support in the form of pipelining the lookup stages.

[064] Although the present invention has been described with an exemplary embodiment, various changes and modifications may be suggested to one skilled in the art. It is intended that the
10 present invention encompass such changes and modifications as fall within the scope of the appended claims.